Top Ten Transhumanist Technologies

by Lifeboat Foundation Scientific Advisory Board member Michael Anissimov.

Overview

Transhumanists advocate the improvement of human capacities through advanced technology. Not just technology as in gadgets you get from Best Buy, but technology in the grander sense of strategies for eliminating disease, providing cheap but high-quality products to the world's poorest, improving quality of life and social interconnectedness, and so on. Technology we don't notice because it's blended in with the fabric of the world, but would immediately take note of its absence if it became unavailable. (Ever tried to travel to another country on foot?) Technology needn't be expensive — indeed, if a technology is truly effective it will pay for itself many times over.

Transhumanists tend to take a longer-than-average view of technological progress, looking not just five or ten years into the future but twenty years, thirty years, and beyond. We realize that the longer you look forward, the more uncertain the predictions get, but one thing is quite certain: if a technology is physically possible and obviously useful, human (or transhuman!) ingenuity will see to it that it gets built eventually.

As we gain ever greater control over the atomic structure of matter, our technological goals become increasingly ambitious, and their payoffs more and more generous. Sometimes new technologies even make us happier in a long-lasting way: the Internet would be a prime example. In the following list I take a look at what I consider the top ten transhumanist technologies.

The List

10. Cryonics

Cryonics is the high-fidelity preservation of the human body, and particularly the brain, after what we would call death, in anticipation of possible future revival. Cryonics is an important transhumanist technology not only because it is already available today, but because the technology is relatively mature — we can reliably stop cells from decaying. In vitrification, the brain is not frozen in the conventional manner but with a cryoprotectant (antifreeze) mixture, which effectively prevents the formation of crystals, causing the water to freeze smoothly, like glass.

Maintenance of a cryo-patient is not difficult — it requires no electricity, but merely the replenishment of liquid nitrogen about every three weeks. As cryonics becomes more popular, this process could become automated and extremely reliable. Further improvements in dewar technology will continue to increase safety and reduce costs. The Cryonics Institute in Michigan, for example, has operated since 1976 without a single mishap.

Financed by the interest of the payout of a life insurance policy (which for people under 40 may cost as little as $100 a year to own), patients can be securely cryopreserved for as long as the cryonics company stays afloat and the dewar stays in one piece. Eventual revival does not require the technology to become available tomorrow, or next year… as long as the liquid nitrogen keeps replenished, you can stay on ice for as long as it takes.

For an existence proof of cryonic revival, there are frogs that can freeze solid and revive later, though reviving a human from freezing would likely require molecular nanotechnology (MNT). When we will be able to revive a cryo-patient will be strongly related to when we develop sophisticated MNT. Once we do develop MNT, the prospect of successful revival is extremely likely — it would involve slowly melting the ice and rebooting the metabolism by kickstarting the appropriate chemical reactions within cells.

9. Virtual Reality

The above image may look like a photo, but it’s actually a screenshot from the game Crysis, a first-person shooter which will be released later this year. Look at screenshots from the game and you’ll see that computer graphics are already beginning to approach photorealism. Sometime in the 2020s, reality simulations will become

so high-resolution and immersive that they'll start to get indistinguishable from the real thing.

Simulations will become the preferred environments for work and play. Pretty soon the main obstacle to truly immersive VR will not be the visuals but the haptics — our sense of touch. To fool our senses into believing haptic technologies are conveying the real thing, the "frame rate" needs to be significantly higher than for visual technologies, a few hundred updates per second rather than a few dozen — which is why development could take another decade or two. But many millions of dollars are currently going into efforts to develop advanced VR.

Clearly, World of Warcraft's eight million subscribers and Second Life's five million subscribers are onto something. At least 1% of all broadband Internet users play in virtual worlds, and this number is increasing rapidly. These worlds typically outclass the real world in terms of customizability, but still have yet to catch up in terms of sensory richness or social fulfillment. But it's only a matter of time.

In the mid-to-late 2020s, I expect full-body, high quality haptic VR suits to be affordable to the average person in developed countries, obtained either from your local Wal-Mart or perhaps printed right out of a desktop nanofactory after payment of a fee. For more on this, here is one scientific paper, "Towards full-body haptic feedback".

8. Gene Therapy/RNA Interference

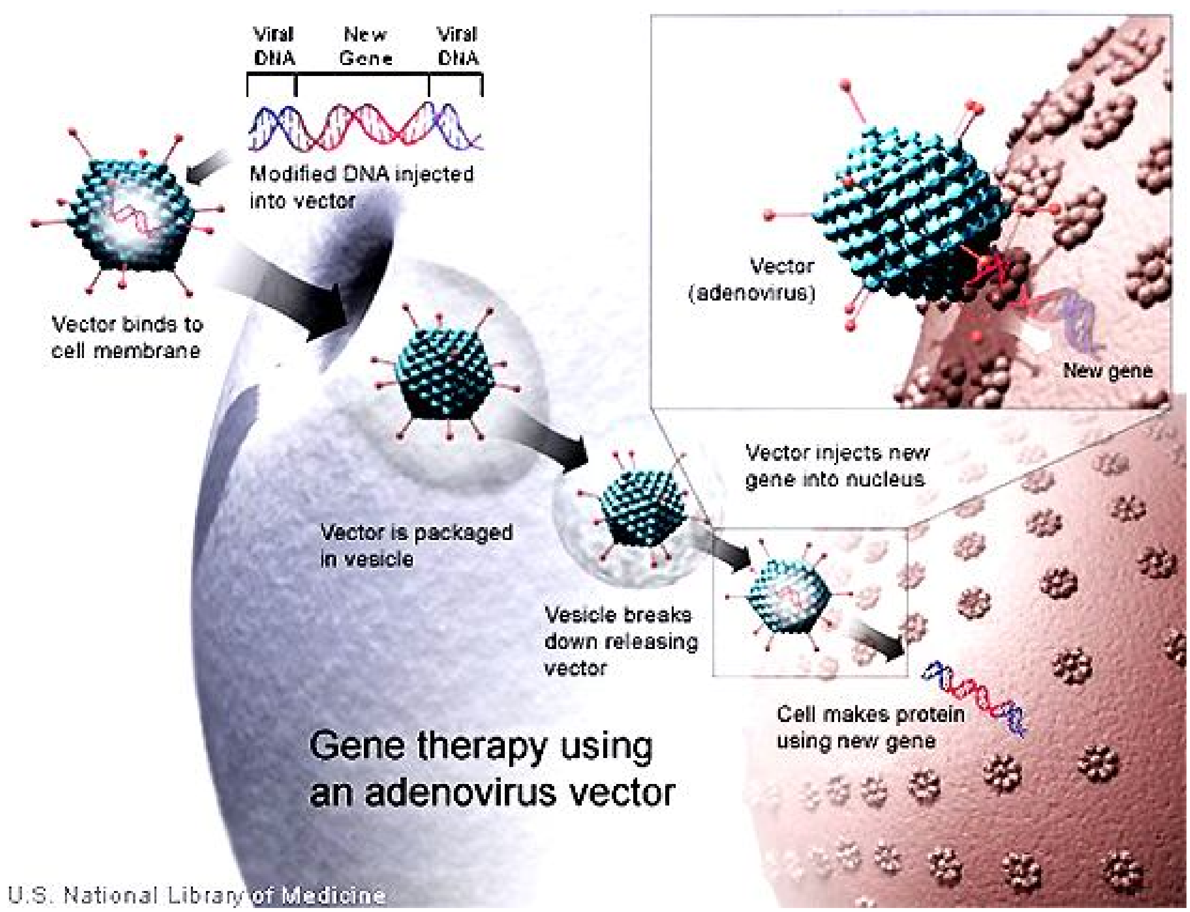

Gene therapy replaces bad genes with good genes, and RNA interference can selectively knock out gene expression. Together, they give us an unprecedented ability to manipulate our own genetic code. By knocking out genes that code for certain metabolic proteins, scientists have been able to make mice that stay slim no matter how much junk food they eat. Lou Gehrig's disease has been cured in mice, and it could only be a few years before we develop a therapy that can cure it for humans too.

Aubrey de Grey's SENS (Strategies for Engineered Negligible Senescence) research program contains various prescriptions for the use of gene therapy. Within a couple decades or so, progress in anti-aging therapies will improve to the point where we are gaining more than an extra year of lifespan per year, reaching so-called "longevity escape velocity" eventually culminating in indefinite lifespans.

Like many transhumanist technologies, gene therapy is really exciting because it's just beginning. No scientist has yet performed gene therapy on germline cells (sexual cells in the gonads) due to the ethical controversy of producing genetic changes which are heritable, but, as with many of these things, it's only a matter of time. Regulations in any given country will only be capable of slowing the overall progress of the field by a few years at most.

The money will go where the research is permitted. In its mature form, gene therapy and genetic engineering will become extremely cheap and powerful, letting humans live comfortably in a wider range of environments and gain immunity to most, if not all diseases. Supercomputers of the future, with thousands or millions of times the crunch power of today's best, will let us simulate the changes in extreme detail before we attempt them with actual human beings. This will make ill side effects quite unlikely for the typical case, much to the dismay of the authors of "genetic engineering turned daddy into a bloodthirsty zombie!" trash novels and films.

7. Space Colonization

Space colonies will become necessary to house the many billions of individuals that will be born in the future as our population continues to expand at a lazy exponential. In his book, *The Millennial Project*, Marshall T. Savage estimates that the Asteroid Belt could hold 7,500 *trillion* people, if thoroughly reshaped into O'Neill colonies. At a typical population growth rate for developed countries at 1% per annum (doubling every 72 years), it would take us 1,440 years to fill that space. Siphoning light gases off Jupiter and Saturn and fusing them into heavier elements for construction of further colonies seems plausible in the longer term as well.

Why expand into space? For many, the answers are blatantly obvious, but the easiest is that the alternatives are limiting the human freedom to reproduce, or mass murder, both of which are morally unacceptable. Population growth is not inherently antithetical to a love of the environment — in fact, by expanding outwards into the cosmos in all directions, we'll be able to seed every star system with every species of plant and animal imaginable. The genetic diversity of the embryonic home planet will seem tiny by comparison.

Space colonization is closely related to transhumanism through the mutual association of futurist philosophy, but also more directly because the embrace of transhumanism will be necessary to colonize space. Human beings aren't designed to live in space. Our physiological issues with it are manifold, from deteriorating muscle mass to uncontrollable flatulence. On the surface of Venus, we would melt, on the surface of Mars, we'd freeze. The only reasonable solution is to upgrade our bodies. Not terraform the cosmos, but cosmosform ourselves.

6. Cybernetics

Can you spot the cyborg in this picture? You're looking right at him! It's Michael Chorost, the man who was born almost deaf but now can hear, thanks to a cochlear implant. Most of the cyborgs in fiction fit certain stereotypes — bermensch wannabes, cyborg assassins, and supercops. But cyborgs already walk among us,

and they look just like normal people.

This trend will continue in the future. Many cyborg upgrades which will become available in the 20s and 30s, such as hearing and vision enhancement, metabolic enhancement, artificial bones, muscles, and organs, and even brain-computer interfaces will be invisible to the casual observer, implanted beneath the skin. Cybernetic features on the surface, such as dermal enhancements or technological actuators like retractable wings, will be carefully camouflaged. No one will want to shock the rest of society by looking like the tin man in public.

The process of cyborgization has already been happening for centuries if not millennia, since the advent of clothing and piercings. For many generations, but especially in the last couple decades, our technological gadgets have been getting smaller, more functional, and more closely integrated with our natural activity.

Recently, Microsoft announced Microsoft Surface, a mouseless, keyboardless form of desktop computing which takes input from finger tracing and hand gestures. The sophistication of biotechnology and the availability of better materials and precision manufacturing will let us make systems so small and effective that even everyday people elect to implant them.

These cybernetic systems will greatly improve our everyday experience, from letting us hear a wider range of ambient sounds, to viewing millions of stars rather than just a few thousand, to making us more resistant to accidents. They will improve the overall economy by enabling us do more work in less time for better pay. In the long term, enhanced humans may get a bigger portion of the economic pie than un-augmented humans, but the pie itself will become so much larger than even the poorest humans of tomorrow will be better off than the wealthiest of today.

Here’s a good cyborg blog I found while doing research for this article, and the Power Jacket, a 4-pound jacket that enhances strength and is used by people recovering from paralysis. For more, see the cybernetics category of my del.ic.ious links, or my top ten list of cybernetic enhancements.

5. Autonomous Self-Replicating Robotics

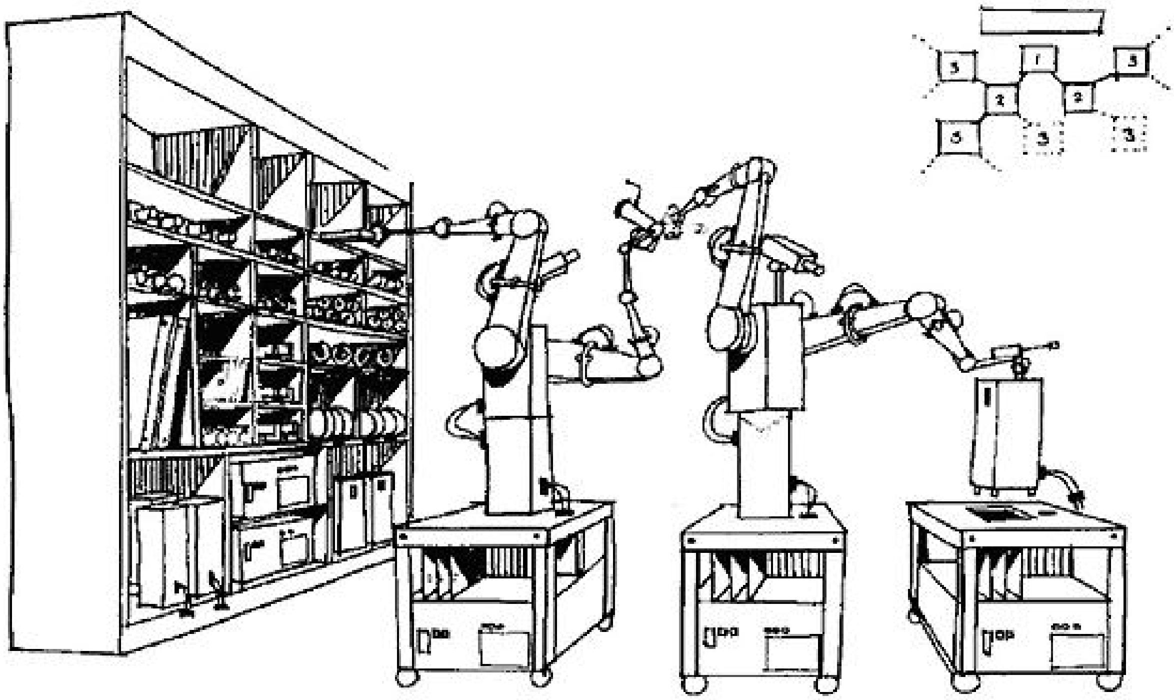

Why do manual labor when the robots can do it for you? Self-replication might be considered the Holy Grail of robotics. A landmark NASA study, “Advanced Automation for Space Missions”, found that robotic self-replication is just a matter of engineering, and that no fundamental theoretical breakthroughs are needed. The study proposed sending a 100-ton package to the Moon, with a self-replication time of 1 year, and letting it self-replicate until the desired level of development is attained.

The design — which was fleshed out in great detail — was based on electric carts running on rails within the factory, “paving machines” that direct sunlight to melt lunar regolith, robotic strip miners for obtaining raw materials, and a solar cell “canopy” for powering it all. After 10 years, over 100,000 tons of lunar factory could be produced autonomously. The factory’s functions could then be hijacked for the benefit of human colonists, used to produce housing, products, and provide large quantities of solar power.

If similar self-replicating systems could be constructed on Earth, there would be little limit to the material plenty they could provide. Self-replicating factories could turn the vast empty badlands of Australia into lush gardens by pumping water from the oceans, self-replicating factories in the high Arctic could melt snow and create gigantic transparent domes suitable for habitation, and submersible automata in the seas could dredge sand from abiotic regions of the ocean floor and process it into gigantic platforms for human colonization.

By opening up such vast new regions of the Earth’s surface, talk of overpopulation and crowding would fall by the wayside for quite a few decades, with people realizing how much space there actually was all along. And once things really do get too crowded here on Earth, we can move to the Moon, Mars, and the asteroid belt, using the power of self-replicating robotics to create rotating space colonies suitable for housing trillions of people.

Self-replicating factories could reduce the costs of material goods close to that of food — the primary expenses would consist of raw materials, energy, and whatever small quantity of human oversight is necessary to keep an eye on the overall structure of things. By utilizing special, man-made “nutrients” for top-level functions (rare or

exotic molecules such as custom-synthesized proteins) and the broadcast architecture — whereby derivative factories must receive affirmations from a central parent factory to continue self-replicating — such factories could be made safe by design.

With such abundance, humanity might actually shift from having a zero-sum perspective on a world to a positive-sum perspective. With medical tools and basic goods in ample supply, no one in the world would need to suffer from poverty or curable disease. The nature of human work would shift from manual drudgery and mind-numbing routine to more creative and personally fulfilling endeavors, like art, music, math, science, literature, and exploration.

For more details on the state of the art in self-replicating machines, see the Wikipedia entry, or the magnum opus on the topic, *Kinematic Self-Replicating Machines*.

4. Molecular Manufacturing

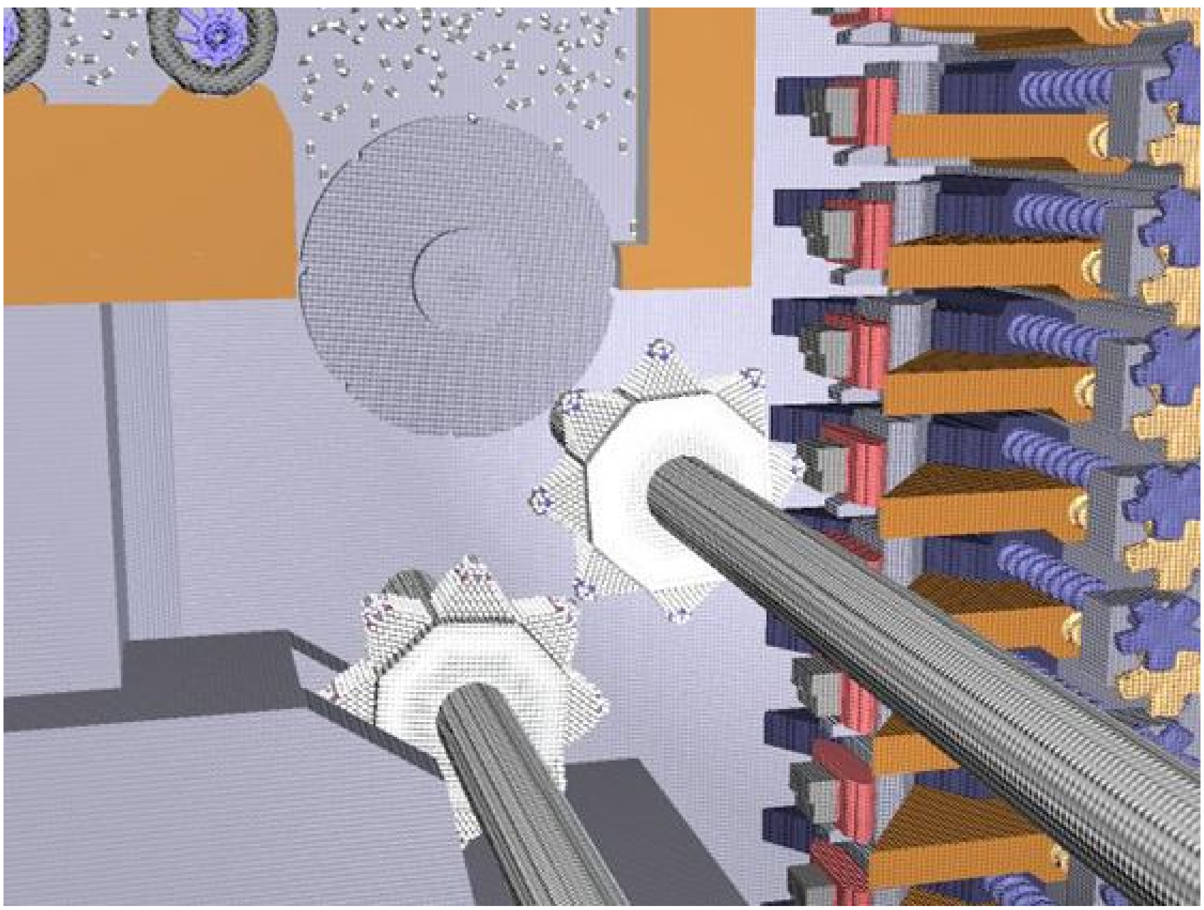

Nanofactory in action. Courtesy Lizard Fire Studios.

If self-replication is the Holy Grail of robotics, then molecular nanotechnology (MNT) is the Holy Grail of manufacturing. Molecular nanotechnology would use massive arrays of nanometer-scale actuators (produced initially through self-replication) to manufacture macroscale products with atomic precision. This concept is known as the nanofactory.

In practical terms, the creation of nanofactories would mean that practically everything could be made out of diamond, motors would become so powerful that a cubic centimeter would provide enough torque to propel a car, medical nanodevices could heal wounds and repair organs without the need for surgery, and air-suspended nanodevices ("utility fog") could be configured to simulate practically any desired object on demand.

On the downside of things, it could become easy to manufacture mite-sized robots with a payload of poison sufficient to kill thousands, or a laptop-sized device capable of separating U-235 from U-238 in a worrisomely simple and rapid fashion, or self-replicating synthetic algae capable of clogging up our oceans with grey goo.

Enabling widespread use of the positive applications while cleanly and completely suppressing the nasty applications is a first-order challenge. Incidentally, you can make a difference right now by donating to the Lifeboat Foundation or Center for Responsible Nanotechnology, two of very few organizations focusing on this area.

To some, molecular nanotechnology sounds like science fiction, and based on the grandiose applications I discussed in the previous paragraph, you can't blame them. But many of the prerequisites of molecular manufacturing have already been demonstrated — "molecular surgery" has been used to snip off and replace individual hydrogen atoms, various functional nanoscale devices have been built, scanning tunneling microscopy has been used to mechanically manipulate individual atoms, and so on. The challenge is to create a nanoscale manipulator arm capable of placing individual atoms with angstrom-level precision, avoiding undesired reactions, and serving as a universal constructor that can build a copy of itself.

There are numerous technical challenges still outstanding, but when these are overcome, manufacturing will be granted the power that nature has had for hundreds of millions of years — the ability to fabricate large objects with molecular precision. The numerous potential applications of the technology to human enhancement are obvious; with molecular manufacturing, we could orchestrate elegant improvements to every single body component, achieving all of the upgrades described on my top ten list, and many more.

3. Megascale Engineering

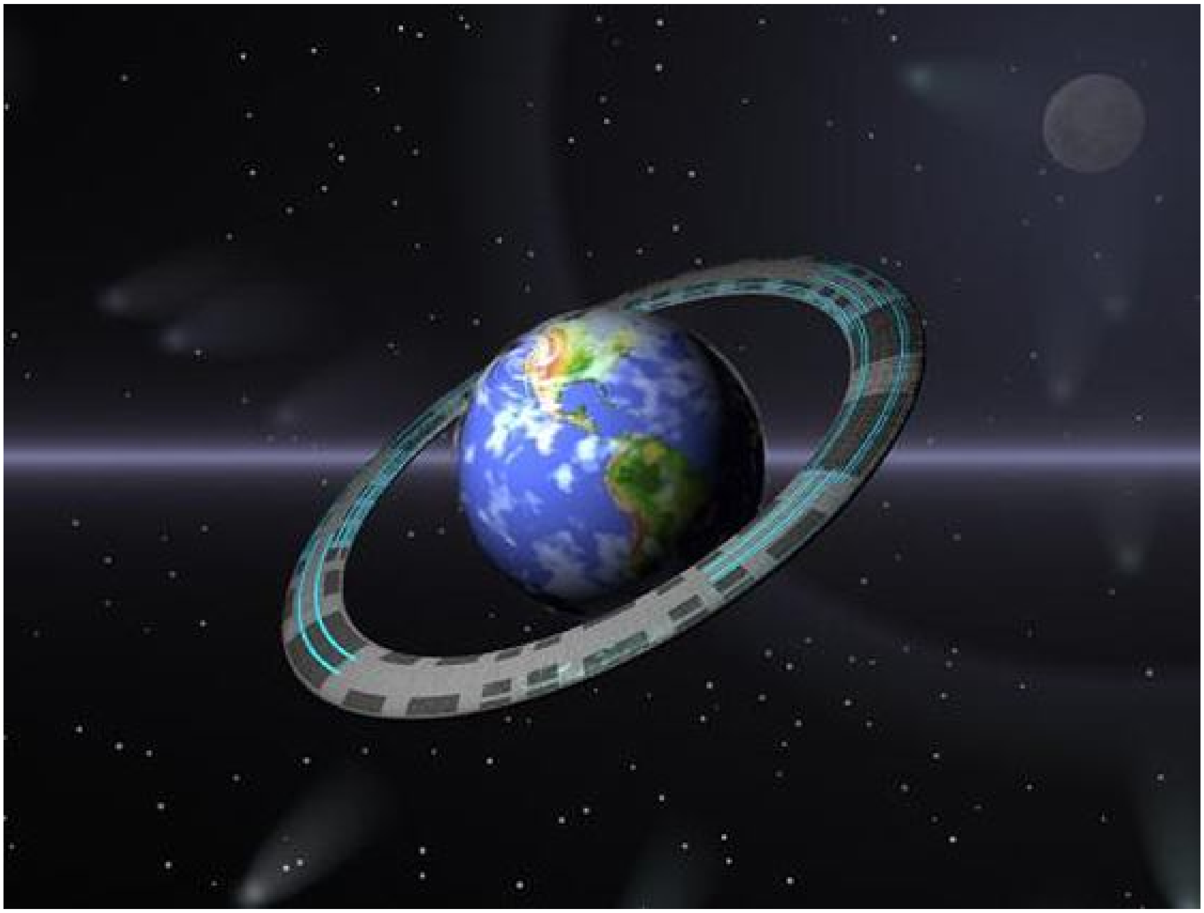

Most people are familiar with megascale engineering because it is seen throughout fiction — the Death Star, for instance. Typically, megascale engineering refers to building structures at least 1,000 km in length in one dimension, such as a space elevator, Globus Cassus, or Dyson sphere. With the self-replicating robotics described above, the production of such large structures could be done largely by autonomous drones, with intelligent agents only managing the highest top-level functions and architecture.

Considering that mankind's long-term future is in space, and that space right now is pretty devoid of any structure useful or habitable to humans, we have a lot of work to do, and if you can make the projects megascale, why not?

Like some of the other items on this list, megascale engineering is only indirectly transhumanist — but is still very relevant to the long-term future of intelligent life. Megascale engineering goes hand-in-hand with the grandiose transhumanist vision: intelligent beings spreading across the cosmos, and eventually shaping the very structure of the universe itself.

The fact that these vast expanses of colonizable space are currently neglected imposes on us a vast opportunity cost — if we hurried up a bit and colonized them, we could give rise to tremendous numbers of people leading worthwhile lives. What experiences would they have, and what stories would they tell? We'll never find out, unless we make it happen.

2. Mind Uploading

Mind uploading, sometimes referred to as nonbiological intelligence, centers around the controversial proposition that cognitive processing can be implemented on substrates other than our current neurons. Considering decades of successful results in neurophysiology, and the recent construction of the world's first brain prosthesis — an artificial copy of the hippocampus — this seems very likely. It appears that our minds are defined more by the information pattern they embody than the particular hardware they are implemented on.

Numerous philosophers of mind have long recognized this, but acceptance among the wider public has been a long time in coming: people don't want to think that they're "just" data structures being implemented as computational automata on biological neurons. But it is hard to think of it any other way: once we dismiss the possibility of an immaterial soul, we must acknowledge the mind as a material pattern implemented in physical configurations, and if other substances aside from our current neurons can meet the requirements for these configurations, then there is no reason why intelligence and consciousness could not exist on another substrate. For a humorous look at this complex philosophical argument, see "They're Made Out of Meat" by Terry Bisson.

If our brains really don't have to be made out of meat, then we can transfer them to other substrates. By incrementally replacing each neuron with a synthetic neuron-equivalent, the whole process could go down painlessly and seamlessly. The transfer could be as slow or as fast as we want: from the information-processing perspective of the brain itself, nothing ever changes. Light still comes in through the eye's lens, hits the retina, is transformed into nerve impulses which travel down the optic nerve, receives further processing in the visual cortex at the back of the brain, the highlights of which are sent to the prefrontal cortex for integration with information from the other senses.

The brain can't tell if it's made out of traditional meat, or accelerated biological neurons, or entirely nonbiological

neuron-equivalents: the computation is the same. Sometimes this notion is also referred to as an application of the Church-Turing thesis.

If entirely synthetic brains are possible, then there's nothing stopping such persons from inhabiting computer networks — not indirectly, sitting in chairs as we currently do, but directly, engaging in computer worlds as a sentient program of tremendous complexity. With molecular manufacturing on hand, reversing the process would be as simple as printing out a hundred or so kilograms of flesh and bone again, complete with memories from the networked experience.

This is probably among the transhumanist visions that most reliably elicits the "yuck!" reaction, but if functionalism is true, then virtual experience will be indistinguishable from physical experience. Not only that, but even more enjoyable, due to the manifold degrees of freedom which would become newly accessible. In a virtual world, there are no laws of physics except those we choose.

For a bit more on simulations, see this primer, and remember: be careful not to generalize from fictional evidence.

1. Artificial General Intelligence (AGI)

As argued in the previous section, functionalism seems likely. If so, then strong AI is possible. Thinking, feeling, imagining, creating, communicating, thoughtful synthetic intelligences with conscious experiences. Whether serial computing is sufficient, or parallel computing is necessary, both are within technological reach, and present-day computing speeds are fast approaching the computing power of the human brain.

In fact, according to many estimates, the fastest present-day supercomputer, Blue Gene/P, has already exceeded it. Blue Gene/P operates continuously at speeds of over a petaflop, which is a million billion operations per

second. For strong AI skeptics, no computer — even one operating at trillions of trillions of trillions of operations per second, is sufficient to implement true intelligence, but to functionalists like myself, such a meat-centric perspective is unjustified.

Distinct from artificial intelligence in general, which has come to refer to any sophisticated software program, artificial general intelligence refers to AIs that display open-ended learning and similar competency levels to human beings.

A handful of researchers are working diligently towards artificial general intelligence, informed by the mathematics of inference and probability theory: Jürgen Schmidhuber, whose "main scientific ambition has been to build an optimal scientist, then retire"; Marcus Hutter, author of the landmark book *Universal Artificial Intelligence*; Ben Goertzel, who recently presented his AI design in a talk to Google; and Eliezer Yudkowsky, who is developing a reflective decision theory from first principles. Whether or not others believe in the feasibility of general AI, these individuals will keep working, and one will eventually succeed.

The way the world would be impacted by the arrival of general AI is too extreme to discuss in much detail here. If raw materials such as sand can be converted into computer chips and then into intelligent minds, eventually the majority of material in the solar system could be made intelligent and conscious. The result would be a "noetic Renaissance": the expansion of intelligence and experience beyond our wildest dreams.

Conversely, if not given empathic values, artificial intelligence could lead to the doom of all. It's up to us to set the initial conditions appropriately: if not, we might not be around to regret it.